







Exhibit 4

U.S. Patent No. 7,784,058 vs. Oracle

Accused Instrumentalities: Oracle products and services using user mode critical system elements as shared libraries, including without limitation Oracle Cloud Infrastructure (“OCI”) and Oracle Kubernetes Engine (“OKE”), and all versions and variations thereof since the issuance of the asserted patent.

Claim 1

Claim 1	Accused Instrumentalities
[1pre] 1. A computing system for executing a plurality of software applications comprising:	To the extent the preamble is limiting, each Accused Instrumentality comprises or constitutes a computing system for executing a plurality of software applications as claimed. <i>See claim limitations below.</i> <i>See also, e.g.:</i>

Claim 1	Accused Instrumentalities
	<div data-bbox="699 212 1094 261"> <h2>Why Choose OKE?</h2> </div> <div data-bbox="699 337 1094 651">  <p>Price-Performance</p> <p>OKE is the lowest cost Kubernetes service amongst all hyperscalers, especially for serverless.</p> </div> <div data-bbox="1129 337 1535 651">  <p>Autoscaling</p> <p>OKE automatically adjusts compute resources based on demand, which can reduce your costs.</p> </div> <div data-bbox="1570 337 1934 651">  <p>Efficiency</p> <p>GPUs can be scarce, but OKE job scheduling makes it easy to maximize resource utilization.</p> </div> <div data-bbox="699 727 1094 1024">  <p>Portability</p> <p>OKE is consistent across clouds and on-premises, enabling portability and avoiding vendor lock-in.</p> </div> <div data-bbox="1129 727 1535 1024">  <p>Simplicity</p> <p>OKE reduces the time and cost needed to manage the complexities of Kubernetes infrastructure.</p> </div> <div data-bbox="1570 727 1934 1024">  <p>Reliability</p> <p>Automatic upgrades and security patching boost reliability for the control plane and worker nodes.</p> </div> <div data-bbox="674 1068 1480 1105"> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/</p> </div>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="693 217 1696 269">Welcome to Oracle Cloud Infrastructure</h2> <p data-bbox="693 321 1944 467">Oracle Cloud Infrastructure (OCI) is a set of complementary cloud services that enable you to build and run a range of applications and services in a highly available hosted environment. OCI provides high-performance compute capabilities (as physical hardware instances) and storage capacity in a flexible overlay virtual network that is securely accessible from your on-premises network.</p> <p data-bbox="674 492 1667 521">https://docs.oracle.com/en-us/iaas/Content/GSG/Concepts/baremetalintro.htm</p> <h3 data-bbox="722 583 1871 695">Existing applications can benefit by migrating to OCI and OKE</h3> <p data-bbox="722 745 1797 781">OKE offers lower total cost of ownership and improved time to market.</p> <p data-bbox="722 846 1572 881">OKE simplifies operations at scale in the following ways:</p> <ul data-bbox="722 971 1871 1344" style="list-style-type: none"> <li data-bbox="722 971 1251 1052">– Lift and shift; there’s no need to rearchitect <li data-bbox="1318 971 1854 1052">– Increase resource utilization and efficiency <li data-bbox="722 1117 1241 1198">– Reduce operations burden with automation <li data-bbox="1318 1117 1871 1198">– Improve agility, flexibility, uptime, and resilience <li data-bbox="722 1263 1178 1344">– Save time on infrastructure management <li data-bbox="1318 1263 1787 1344">– Reduce compliance risk and enhance security <p data-bbox="674 1393 1675 1422">https://www.oracle.com/cloud/cloud-native/kubernetes-engine/#app-migration</p>

Claim 1	Accused Instrumentalities
	<p>What is OCI Kubernetes Engine (OKE)?</p> <p>Oracle Cloud Infrastructure Kubernetes Engine (OKE) is a managed Kubernetes service that simplifies the development, deployment, and operation of containerized workloads at scale. OKE enables you to quickly create, manage, and consume Kubernetes clusters that leverage underlying OCI compute, networking, and storage services.</p> <p>When should I use OKE?</p> <p>You should use OKE when you want to leverage Kubernetes to deploy and manage your Kubernetes-based container applications. It allows you to combine the production-grade container orchestration of standard upstream Kubernetes with the control, security, and high, predictable performance of OCI.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <p>How does OKE provide resiliency?</p> <p>When you create an OKE cluster, OKE automatically creates and manages multiple Kubernetes control plane nodes spread across fault domains and availability domains (logical data centers). This is done to help ensure that the managed Kubernetes control plane is highly available. Control plane operations, such as upgrading to newer versions of Kubernetes, can be performed without service interruptions. Additionally, when you provision worker nodes, you can use a placement configuration to control the fault domain and availability domain where they are created. Nodes will automatically come online with labels, which you can use to schedule your workloads so they are robust and highly available.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <p>Can I deploy private Kubernetes clusters?</p> <p>Yes; with OKE, your Kubernetes clusters are integrated in your VCN. Your cluster worker nodes, load balancers, and the Kubernetes API endpoint are part of a private or public subnet of your VCN. Regular VCN routing and firewall rules control access to the Kubernetes API endpoint, making it accessible from a corporate network only, through a bastion host, or by specific platform services.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p>

Claim 1	Accused Instrumentalities
	<p>When should I use virtual nodes, managed nodes, or self-managed nodes?</p> <ul style="list-style-type: none"> Virtual nodes Virtual nodes offer a serverless Kubernetes experience. This option is ideal if you'd rather focus on your application and avoid managing the underlying infrastructure. Virtual nodes relieve you of management-related tasks such as scaling, upgrading, patching, troubleshooting, and provisioning worker nodes. Managed nodes Managed nodes are a good choice for general purpose workloads. They offer an extensive list of customizable configuration options that have been tested by the OKE service. Unlike fully managed virtual nodes, you share the management of worker nodes with OCI. OKE simplifies the management process through features such as on-demand cycling to automate worker node updates, cluster self-healing upon failure detection, autoscaling, and more. Self-managed nodes Self-managed nodes offer access to the underlying infrastructure, configuration options, and compute shapes that aren't currently available to managed nodes. This includes access to specialized infrastructure, such as RDMA-enabled bare metal cluster networks or confidential compute shapes. This advanced control makes self-managed nodes ideal for specialized use cases that aren't supported with managed nodes. Note that with self-managed nodes, you are fully responsible for managing the worker nodes—without the automated features provided by managed or virtual nodes. <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <p>What are the storage options for virtual nodes?</p> <p>OKE virtual nodes do not yet have persistent storage capabilities. However, there are plans to introduce support for attaching persistent volumes backed by OCI Block Storage and OCI File Storage. If your Kubernetes application requires persistent storage, it's advisable to use OKE managed nodes.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p>

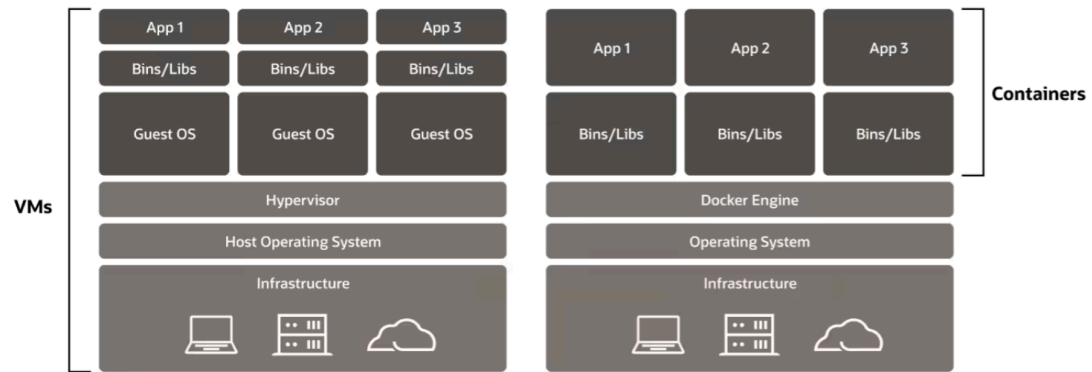
Claim 1	Accused Instrumentalities
	<h2 data-bbox="709 215 1581 267">Supported Images for Managed Nodes</h2> <p data-bbox="709 313 1959 380">Kubernetes Engine supports the provisioning of worker nodes (managed nodes only) using some, but not all, of the latest Oracle Linux images provided by Oracle Cloud Infrastructure.</p> <p data-bbox="709 425 1602 451">You can use these Oracle Linux images when provisioning managed nodes:</p> <ul data-bbox="737 496 972 638" style="list-style-type: none"> • OKE Images • Platform Images • Custom Images <p data-bbox="674 672 1818 703">https://docs.oracle.com/en-us/iaas/Content/ContEng/Reference/contengimagesshapes.htm</p> <h2 data-bbox="684 743 1014 787">What is Docker?</h2> <hr data-bbox="684 824 749 834"/> <p data-bbox="684 896 1944 1052">A Docker container is a packaging format that packages all the code and dependencies of an application in a standard format that allows it to run quickly and reliably across computing environments. A Docker container is a popular lightweight, standalone, executable container that includes everything needed to run an application, including libraries, system tools, code, and runtime. Docker is also a software platform that allows developers to build, test, and deploy containerized applications quickly.</p> <p data-bbox="684 1086 1959 1177">Containers as a Service (CaaS) or Container Services are managed cloud services that manage the lifecycle of containers. Container services help orchestrate (start, stop, scale) the runtime of containers. Using container services, you can simplify, automate, and accelerate your application development and deployment lifecycle.</p> <p data-bbox="684 1211 1902 1302">Docker and Container Services have seen rapid adoption and have been a tremendous success over the last several years. From an almost unknown and rather technical open source technology in 2013, Docker has evolved into a standardized runtime environment now officially supported for many Oracle enterprise products.</p> <p data-bbox="674 1330 1707 1360">https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>Docker versus Kubernetes</p> <p>Linux containers have existed since 2008, but they were not well known until the emergence of Docker containers in 2013. With the onset of Docker containers, came the explosion of interest in developing and deploying containerized applications. As the number of containerized applications grew to span hundreds of containers deployed across multiple servers, operating them became more complex. How do you coordinate, scale, manage, and schedule hundreds of containers? This is where Kubernetes can help. Kubernetes is an open source orchestration system that allows you to run your Docker containers and workloads. It helps you manage the operating complexities when moving to scale multiple containers deployed across multiple servers. The Kubernetes engine automatically orchestrates the container lifecycle, distributing the application containers across the hosting infrastructure. Kubernetes can quickly scale resources up or down, depending on the demand. It continually provisions, schedules, deletes, and monitors the health of the containers.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1**Accused Instrumentalities****Docker Basics**

The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.

A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.

**Virtual Machines**

- Each virtual machine (VM) includes the app, the necessary binaries and libraries and an **entire guest operating system**

Containers

- Containers include the app and all of its dependencies, but **share the kernel** with other containers.
- Run as an isolated process in userspace on the host operating system.
- **Not** tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud.

<https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/>


Claim 1	Accused Instrumentalities
	<p>Container Cloud Services</p> <p>The first part of this article explained some important Docker concepts. However, in a production environment it is not enough to simply run an application in a Docker container.</p> <p>To setup and operate a production environment requires hardware to run the containers. Software such as Docker, along with repositories and cluster managers, must be installed, upgraded and patched. If several Docker containers communicate across hosts, a network must be created. Clustered containers should be restarted if they fail. In addition, a set of containers linked to each other should be deployable as easily as a single logical application instance. An example of this could be a load balancer, a few web servers, some Oracle WebLogic Server instances with an admin server, a managed server, and a database. To manage containerized applications at scale, requires a container orchestration system like Kubernetes or Docker Swarm. Deploying, managing, and operating orchestration systems like Kubernetes can be challenging and time-consuming.</p> <p>To make it easier and more efficient for developers to create containerized applications, cloud providers offer Container Cloud Services or Containers as a Service (CaaS). Container Cloud Services help developers and operations teams streamline and manage the lifecycle of containers in an automated fashion. These orchestration services, typically built using Kubernetes, make it easier for DevOps teams to manage and operate containerized applications at scale. Oracle Cloud Infrastructure Kubernetes Engine and Azure Kubernetes Service are two examples of popular container orchestration managed cloud services.</p> <p>Oracle Cloud Infrastructure Kubernetes Engine is a fully managed, scalable, and highly available service that you can use to deploy your containerized applications in the cloud. Use Kubernetes Engine (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud native applications.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
<p>[1a] a) a processor;</p>	<p>Each Accused Instrumentality comprises a processor.</p> <p>For example, each node/host contains at least one CPU.</p> <p><i>See, e.g.:</i></p> <p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="745 690 1820 1055"> </div> <div data-bbox="835 1092 982 1114"> <p>Virtual Machines</p> </div> <ul data-bbox="835 1127 1197 1182" style="list-style-type: none"> • Each virtual machine (VM) includes the app, the necessary binaries and libraries and an <u>entire guest operating system</u> <div data-bbox="1278 1092 1373 1114"> <p>Containers</p> </div> <ul data-bbox="1278 1127 1680 1266" style="list-style-type: none"> • Containers include the app and all of its dependencies, but <u>share the kernel</u> with other containers. • Run as an isolated process in userspace on the host operating system. • <u>Not</u> tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>When should I use virtual nodes, managed nodes, or self-managed nodes?</p> <ul style="list-style-type: none"> Virtual nodes Virtual nodes offer a serverless Kubernetes experience. This option is ideal if you'd rather focus on your application and avoid managing the underlying infrastructure. Virtual nodes relieve you of management-related tasks such as scaling, upgrading, patching, troubleshooting, and provisioning worker nodes. Managed nodes Managed nodes are a good choice for general purpose workloads. They offer an extensive list of customizable configuration options that have been tested by the OKE service. Unlike fully managed virtual nodes, you share the management of worker nodes with OCI. OKE simplifies the management process through features such as on-demand cycling to automate worker node updates, cluster self-healing upon failure detection, autoscaling, and more. Self-managed nodes Self-managed nodes offer access to the underlying infrastructure, configuration options, and compute shapes that aren't currently available to managed nodes. This includes access to specialized infrastructure, such as RDMA-enabled bare metal cluster networks or confidential compute shapes. This advanced control makes self-managed nodes ideal for specialized use cases that aren't supported with managed nodes. Note that with self-managed nodes, you are fully responsible for managing the worker nodes—without the automated features provided by managed or virtual nodes. <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <h2>Welcome to Oracle Cloud Infrastructure</h2> <p>Oracle Cloud Infrastructure (OCI) is a set of complementary cloud services that enable you to build and run a range of applications and services in a highly available hosted environment. OCI provides high-performance compute capabilities (as physical hardware instances) and storage capacity in a flexible overlay virtual network that is securely accessible from your on-premises network.</p> <p>https://docs.oracle.com/en-us/iaas/Content/GSG/Concepts/baremetalintro.htm</p>

Claim 1	Accused Instrumentalities
<p>[1b] b) an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor; and,</p>	<p>Each Accused Instrumentality comprises an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor.</p> <p>For example, the OSCSEs include kernel-mode functions similar to the functionalities provided by user-space libraries such as glibc. These are implemented in kernel-space to handle tasks such as (without limitation) memory management (kmalloc(), kfree(), etc.) at kernel level.</p> <p><i>See, e.g.:</i></p> <p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="745 795 1820 1166"> </div> <div data-bbox="835 1203 982 1222"> <p>Virtual Machines</p> </div> <div data-bbox="835 1237 1197 1289"> <ul style="list-style-type: none"> Each virtual machine (VM) includes the app, the necessary binaries and libraries and an <u>entire guest operating system</u> </div> <div data-bbox="1278 1203 1373 1222"> <p>Containers</p> </div> <div data-bbox="1278 1237 1682 1373"> <ul style="list-style-type: none"> Containers include the app and all of its dependencies, but <u>share the kernel</u> with other containers. Run as an isolated process in userspace on the host operating system. <u>Not</u> tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. </div> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

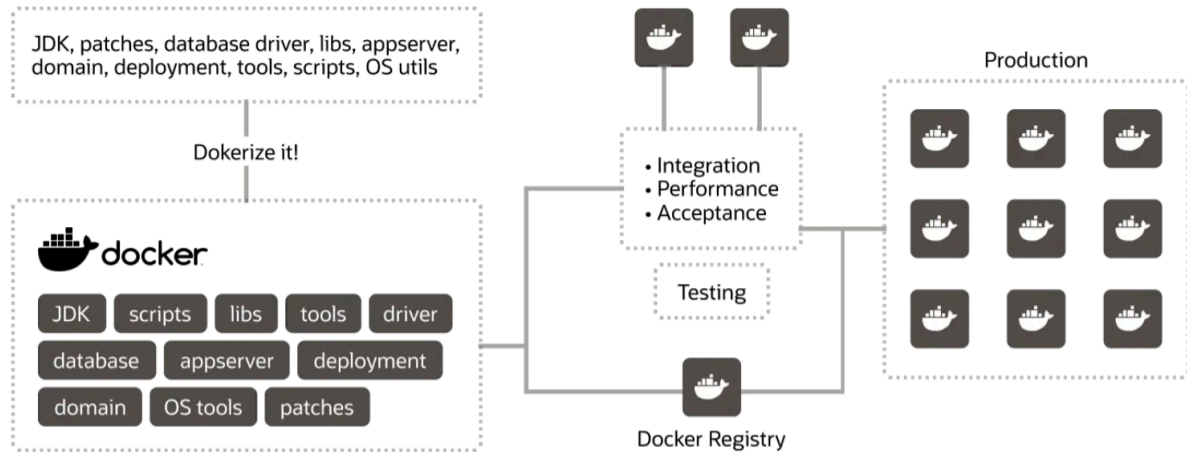
Claim 1	Accused Instrumentalities
	<p>Kernel mode</p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p>https://www.techtarget.com/searchdatacenter/definition/kernel</p> <p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>

Claim 1	Accused Instrumentalities
<p>[1c] c) a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode and</p>	<p>Each Accused Instrumentality comprises a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode.</p> <p>For example, the shared library with SLCSEs include the runtime environment, system tools, and dependencies, such as the glibc library and other libraries that replicate OSCSEs, included in the container image (including without limitation in a base image that is included within the container image).</p> <p><i>See, e.g.:</i></p>  <p>Figure 2</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1**Accused Instrumentalities****Docker Image—Development to Production**

Creating a Docker image with all of its dependencies solves the "but it worked for me on my development machine" problem. The key idea is that a Docker image is created automatically by a build pipeline from a source-code repository like Git and initially tested in a development environment. This immutable image will then be stored in a Docker registry.

As shown in the Figure 4, the same image will be used for further load tests, integration tests, acceptance tests, and more. In every environment, the same image will be used. Small but necessary environmentally specific differences, such as a JDBC URL for a production database, can be fed into the container as environment variables or files.

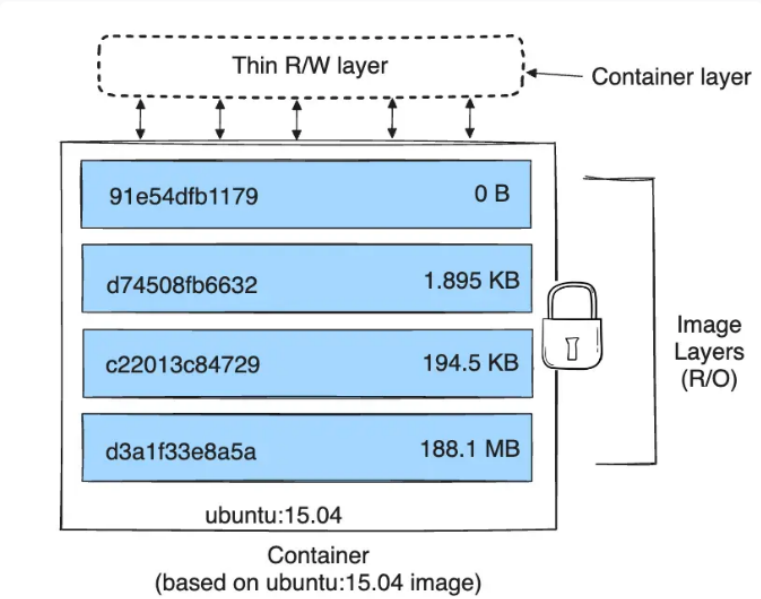


<https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/>

Claim 1	Accused Instrumentalities
	<p>Container images</p> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p> <p>Container image files are complete, static and executable versions of an application or service and differ from one technology to another. Docker images are made up of multiple layers, which start with a base image that includes all of the dependencies needed to execute code in a container. Each image has a readable/writable layer on top of static unchanging layers. Because each container has its own specific container layer that customizes that specific container, underlying image layers can be saved and reused in multiple containers. An Open Container Initiative (OCI)</p> <p>https://www.techtarget.com/searchitoperations/definition/container-containerization-or-container-based-virtualization</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="682 212 1312 277">About storage drivers</h2> <p data-bbox="682 326 1913 448">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 518 1604 574">Storage drivers versus Docker volumes</h2> <p data-bbox="682 612 1953 878">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="682 927 1944 1052">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1381 976 1575 1003" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="674 1084 1266 1117"><a data-bbox="674 1084 1266 1117" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="699 207 1121 264"><h2>Images and layers</h2></div> <div data-bbox="699 302 1860 378"><p>A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p></div> <div data-bbox="699 418 1946 784"><pre># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre></div> <div data-bbox="699 824 1940 1133"><p>This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p></div> <div data-bbox="674 1154 1266 1187"><p>https://docs.docker.com/storage/storagedriver/</p></div>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer structure of a Docker container. At the bottom is a box labeled 'Container (based on ubuntu:15.04 image)'. Inside this box is a stack of four blue rectangular blocks representing image layers. From bottom to top, the layers are labeled with their IDs and sizes: 'd3a1f33e8a5a' (188.1 MB), 'c22013c84729' (194.5 KB), 'd74508fb6632' (1.895 KB), and '91e54dfb1179' (0 B). To the right of this stack is a bracket labeled 'Image Layers (R/O)' with a padlock icon, indicating they are read-only. Above the stack is a dashed box labeled 'Thin R/W layer'. An arrow points from the text 'Container layer' to this dashed box. Vertical double-headed arrows connect the 'Thin R/W layer' to each of the four image layers below it.</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="690 220 959 282"><h2>Volumes</h2></div> <div data-bbox="690 337 1944 467"><p>Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p></div> <div data-bbox="672 488 1348 521"><p>https://kubernetes.io/docs/concepts/storage/volumes/</p></div> <div data-bbox="695 571 1266 620"><h2>Container environment</h2></div> <div data-bbox="695 659 1514 724"><p>The Kubernetes Container environment provides several important resources to Containers:</p></div> <div data-bbox="735 763 1488 922"><ul style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster.</div> <div data-bbox="672 956 1568 989"><p>https://kubernetes.io/docs/concepts/containers/container-environment/</p></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="695 212 917 277"><h2>Images</h2></div> <div data-bbox="695 306 1562 461"><p>A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p></div> <div data-bbox="695 496 1570 570"><p>You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p></div> <div data-bbox="672 596 1369 630"><p>https://kubernetes.io/docs/concepts/containers/images/</p></div> <div data-bbox="690 670 959 729"><h2>Volumes</h2></div> <div data-bbox="690 769 1570 1214"><p>On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p></div> <div data-bbox="672 1240 1348 1274"><p>https://kubernetes.io/docs/concepts/storage/volumes/</p></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="711 219 1337 276"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="711 339 1222 386"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="711 433 1944 506"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="711 542 1944 615"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="669 644 1520 714"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="688 212 869 253"><h3>Overview</h3></div> <div data-bbox="688 306 1932 548"><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p></div> <div data-bbox="688 592 1948 971"><p>The diagram illustrates the components of a container image. On the left, a code block for a <code>HelloWorld</code> class is shown. An arrow points from this code to a cylinder labeled 'layer', which contains the paths <code>/bin/java</code>, <code>/opt/app.jar</code>, and <code>/lib/libc</code>. To the right of the layer is a plus sign, followed by a document icon labeled 'image index' containing a JSON snippet for 'manifests'. Another plus sign follows, then another document icon labeled 'config' containing a JSON snippet for 'config'. Each icon has a small 'Ci' logo in the top left corner.</p></div> <div data-bbox="674 998 1518 1068"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="690 207 1337 267"><h2>OCI Image Configuration</h2></div> <div data-bbox="690 321 1955 483"><p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p></div> <div data-bbox="690 521 1698 553"><p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p></div> <div data-bbox="672 586 1543 656"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p></div>

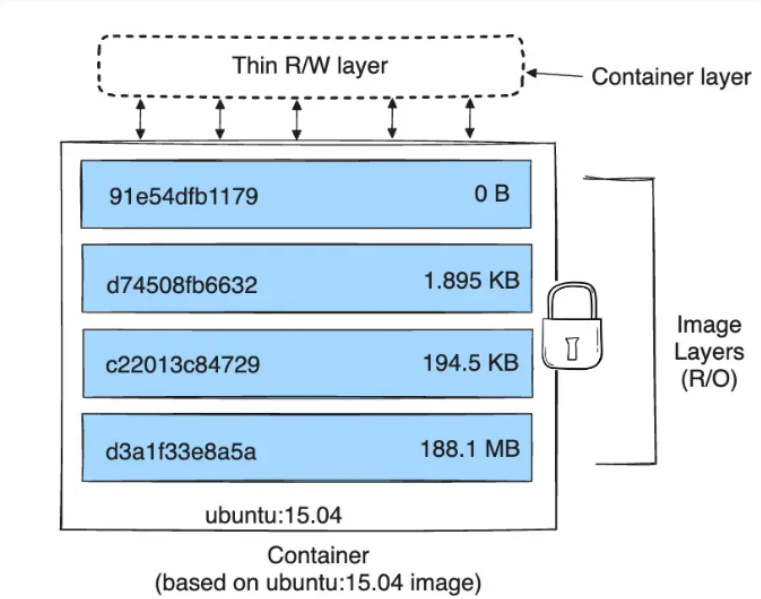
Claim 1	Accused Instrumentalities
	<p data-bbox="701 215 789 253">Layer</p> <ul data-bbox="730 289 1955 634" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="701 686 898 724">Image JSON</p> <ul data-bbox="730 760 1955 1105" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="674 1133 1541 1203"> https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md </p>

Claim 1	Accused Instrumentalities
	<p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>

Claim 1	Accused Instrumentalities
<p>[1d] i) wherein some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications,</p>	<p>In each Accused Instrumentality, some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications.</p> <p>For example, a base image serves as a self-contained unit that encompasses all the necessary components for an application to run, including the application code, runtime environment, system tools, and dependencies (i.e., SLCSEs). The images are based on existing Linux distributions, such as Debian and Ubuntu, including essential system elements (i.e., functional replicas of OSCSEs). Each container image is based on a specific base image, which contains the application code, and dependencies, including system libraries or shared library critical system elements (SLCSEs). The base image forms a part of the container image according to the “layer” model described in the documentation below. When the container runs the image, it creates a runtime instance of that container image. In turn, when one or more applications executes within the container runtime environment, it dynamically links to the SLCSEs stored in the runtime environment, which thereby become a part of the application(s).</p> <p><i>See, e.g.:</i></p> <h2 data-bbox="716 857 1104 914">Container images</h2> <p data-bbox="716 950 1457 1101">A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p data-bbox="674 1125 1268 1156">https://kubernetes.io/docs/concepts/containers/</p>

Claim 1	Accused Instrumentalities
	<p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p> <h2 data-bbox="682 446 1312 511">About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 755 1606 812">Storage drivers versus Docker volumes</h2> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the volumes section to learn how to use volumes to persist data and improve performance.</p> <p>https://docs.docker.com/storage/storagedriver/</p>


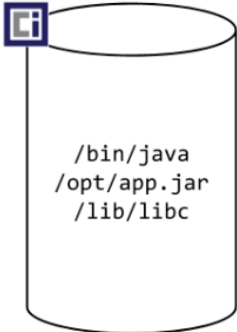

Claim 1	Accused Instrumentalities
	<div data-bbox="699 207 1121 264"><h2>Images and layers</h2></div> <div data-bbox="699 300 1860 378"><p>A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p></div> <div data-bbox="699 418 1946 784"><pre># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre></div> <div data-bbox="699 824 1940 1133"><p>This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p></div> <div data-bbox="674 1154 1266 1187"><p>https://docs.docker.com/storage/storagedriver/</p></div>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer structure of a Docker container. At the bottom is the base image, labeled 'ubuntu:15.04'. Above it is a stack of four image layers, each represented by a blue box with a hash and a size: '91e54dfb1179' (0 B), 'd74508fb6632' (1.895 KB), 'c22013c84729' (194.5 KB), and 'd3a1f33e8a5a' (188.1 MB). A padlock icon is shown next to the stack, labeled 'Image Layers (R/O)'. Above the stack is a dashed box labeled 'Thin R/W layer', with an arrow pointing to it from the label 'Container layer'. Arrows indicate the relationship between the container layer and the image layers.</p> <p>Container (based on ubuntu:15.04 image)</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="690 220 959 282"><h2>Volumes</h2></div> <div data-bbox="690 337 1944 467"><p>Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p></div> <div data-bbox="672 490 1348 522"><p>https://kubernetes.io/docs/concepts/storage/volumes/</p></div> <div data-bbox="697 571 1266 620"><h2>Container environment</h2></div> <div data-bbox="697 659 1514 724"><p>The Kubernetes Container environment provides several important resources to Containers:</p></div> <div data-bbox="735 763 1488 922"><ul style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster.</div> <div data-bbox="672 958 1566 990"><p>https://kubernetes.io/docs/concepts/containers/container-environment/</p></div>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="699 215 915 277">Images</h2> <p data-bbox="699 310 1562 461">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="699 500 1568 570">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="674 597 1367 630">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="695 675 957 737">Volumes</h2> <p data-bbox="695 773 1570 1214">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="674 1242 1346 1274">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<div>Open Container Initiative</div> <div>Image Format Specification</div> <p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<div><h3>Overview</h3><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p><div><div><pre>public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } }</pre></div><div>→</div><div><div><div>/bin/java /opt/app.jar /lib/libc</div></div><div>+</div><div><div><div>{ "manifests": { "platform": { "os": "linux", ... } } }</div></div><div>+</div><div><div><div>{ ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } }</div></div></div><div><div>layer</div><div>image index</div><div>config</div></div></div><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div></div></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="690 207 1337 267"><h2>OCI Image Configuration</h2></div> <div data-bbox="690 321 1955 483"><p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p><p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p></div>

Claim 1	Accused Instrumentalities
	<p data-bbox="701 215 789 253">Layer</p> <ul data-bbox="730 289 1955 634" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="701 686 898 724">Image JSON</p> <ul data-bbox="730 760 1955 1105" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="674 1133 1541 1203"> https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md </p> <p data-bbox="674 1230 1308 1300"> Containers only have access to resources that are defined in the image, https://www.hpe.com/us/en/what-is/docker.html </p>

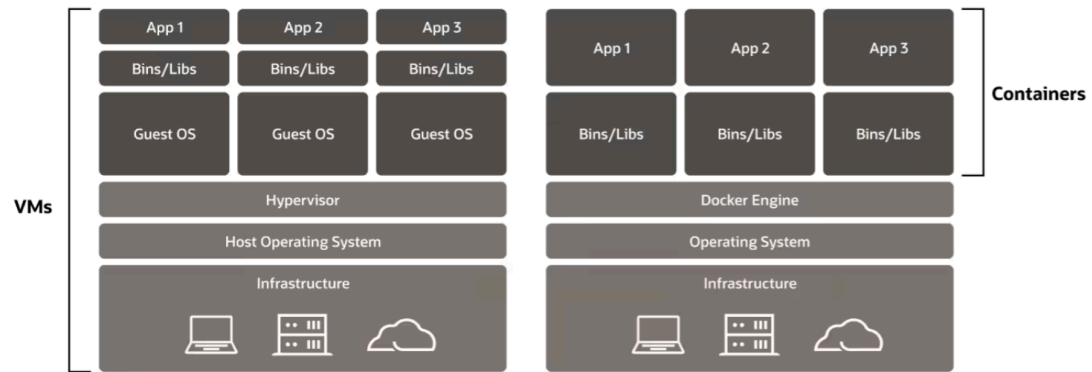
Claim 1	Accused Instrumentalities
	<p>DESCRIPTION top</p> <p>The programs ld.so and ld-linux.so* find and load the shared objects (shared libraries) needed by a program, prepare the program to run, and then run it.</p> <p>https://man7.org/linux/man-pages/man8/ld.so.8.html</p>
<p>[1e] ii) wherein an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function, and</p>	<p>In each Accused Instrumentality, an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function.</p> <p>When a Docker or Kubernetes image is used to create a container in the Accused Instrumentalities, it creates a separate and isolated instance of a runtime environment which is independent of other containers running on the same host. Each container has its own instance of base images and its own data. The containers run in isolation, ensuring that the SLCSEs stored in the shared library are accessible to the software applications running in their respective containers. The image includes essential system files, libraries, and dependencies required to run the software application within the container. The containers can share common dependencies and components using layered images. This means that multiple containers utilize the same base image to create an instance. When an instance of SLCSE is provided from the base image (i.e., from the shared library) to an individual container including application software, it is not shared with other containers. This ensures that each container has its own isolated context. Docker or Kubernetes containers in the Accused Instrumentalities can share common dependencies and components using layered images. This means that multiple containers can utilize the same base image. Therefore, each container, containing the application software running under the operating system, utilizes a unique instance of the corresponding critical system element to execute the respective application software for performing a same function.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p>Cgroups and Namespaces History</p> <p>The underlying Linux kernel features that Docker uses are cgroups and namespaces. In 2008 cgroups were introduced to the Linux kernel based on work previously done by Google developers¹. Cgroups limit and account for the resource usage of a set of operating system processes.</p> <p>The Linux kernel uses namespace to isolate the system resources of processes from each other. The first namespace, i.e. the mount namespace, was introduced as early as 2002.²</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1**Accused Instrumentalities****Docker Basics**

The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.

A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.

**Virtual Machines**

- Each virtual machine (VM) includes the app, the necessary binaries and libraries and an **entire guest operating system**

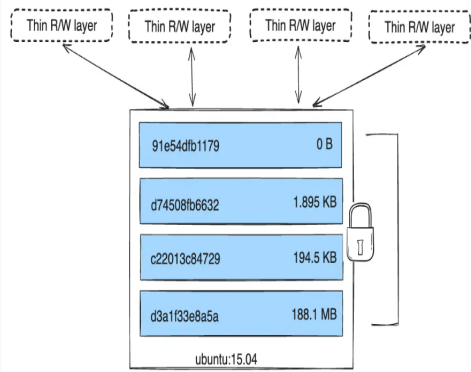
Containers

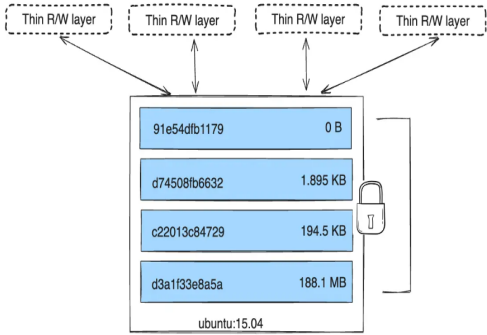
- Containers include the app and all of its dependencies, but **share the kernel** with other containers.
- Run as an isolated process in userspace on the host operating system.
- **Not** tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud.

<https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="709 212 1770 269">Setting Up Storage for Kubernetes Clusters</h2> <p data-bbox="709 310 1955 456"><i>Find out how to define and apply persistent volume claims to clusters you've created using Kubernetes Engine (OKE). With Oracle Cloud Infrastructure as the underlying IaaS provider, you can provision persistent volume claims by attaching volumes from the Block Volume service or by mounting file systems from the File Storage service.</i></p> <p data-bbox="709 500 1917 605">Container storage via a container's root file system is ephemeral, and can disappear upon container deletion and creation. To provide a durable location to prevent data from being lost, you can create and use persistent volumes to store data outside of containers.</p> <p data-bbox="709 649 1866 716">A persistent volume offers persistent storage that enables your data to remain intact, regardless of whether the containers to which the storage is connected are terminated.</p> <p data-bbox="709 760 1934 826">A persistent volume claim (PVC) is a request for storage, which is met by binding the PVC to a persistent volume (PV). A PVC provides an abstraction layer to the underlying storage.</p> <p data-bbox="709 870 1614 899">With Oracle Cloud Infrastructure, you can provision persistent volume claims:</p> <ul data-bbox="735 943 1927 1299" style="list-style-type: none"> • By attaching volumes from the Oracle Cloud Infrastructure Block Volume service. The volumes are connected to clusters created by Kubernetes Engine using CSI (Container Storage Interface) or FlexVolume volume plugins deployed on the clusters. Oracle recommends the CSI volume plugin since the upstream Kubernetes project deprecates the FlexVolume volume plugin in Kubernetes version 1.23. See Provisioning PVCs on the Block Volume Service. • By mounting file systems in the Oracle Cloud Infrastructure File Storage service. The File Storage service file systems are mounted inside containers running on clusters created by Kubernetes Engine using a CSI (Container Storage Interface) volume plugin deployed on the clusters. See Provisioning PVCs on the File Storage Service. <p data-bbox="674 1312 1633 1380">https://docs.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingpersistentvolumeclaim.htm</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="688 212 869 251">Overview</div> <div data-bbox="688 306 1932 548"><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p></div> <div data-bbox="688 592 1948 971"><p>The diagram illustrates the components of an OCI image. On the left, a code block for a 'HelloWorld' class is shown. An arrow points from this code to a cylinder labeled 'layer' containing the paths '/bin/java', '/opt/app.jar', and '/lib/libc'. To the right of the layer is a plus sign, followed by a document icon labeled 'image index' containing a JSON snippet with 'manifests' and 'platform' fields. Another plus sign follows, then a document icon labeled 'config' containing a JSON snippet with 'config' and 'Cmd' fields. Each icon has a small 'Ci' logo in the top left corner.</p></div> <div data-bbox="674 998 1518 1068"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p>
<p>[1f] iii) wherein a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p>	<p>In each Accused Instrumentality, a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p> <p>For example, in Docker or Kubernetes containers used within the Accused Instrumentalities, each container operates independently, and a base image includes essential system files, libraries, and dependencies (i.e., SLCSEs) required to run the software application within the container. Based on information and belief, each element, such as system files, libraries, and dependencies (i.e., SLCSE) is associated with an execution of a predetermined function related to the application. When an image is used to create a container in the Accused Instrumentality, an instance of the SLCSE is provided to a</p>

Claim 1	Accused Instrumentalities
	<p>software application. Therefore, different instances of the SLCSE are provided to different applications for performing a same function, simultaneously.</p> <p><i>See, e.g.:</i></p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p> <p>A container is a runnable instance of an image. You can create, start, stop, move, or delete a container using the Docker API or CLI. You can connect a container to one or more networks, attach storage to it, or even create a new image based on its current state.</p> <p>https://docs.docker.com/get-started/overview/</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p>